

RECEIVED
CENTRAL FAX CENTER

OCT 10 2006

REMARKS

Reconsideration of the present application is respectfully requested. No claims have been amended, canceled or added in this amendment.

Claims 1, 5-7, 11-13, 15-17 and 22-24 stand rejected under 35 U.S.C. § 103(a) based on U.S. Patent no. 6,308,282 of Galipeau et al. ("Galipeau") in view of U.S. Patent Application Publication no. 2004/0230756 of Achiwa et al. ("Achiwa"). Claims 4, 10, 14, 19 and 21 stand rejected under 35 U.S.C. § 103(a) based on Galipeau in view of U.S. Patent Application Publication no. 2003/0014433 of Teloh.

Before discussing the cited art, a brief overview of the invention may be helpful. The present invention relates to a data mirroring system that includes a source storage server which has a persistent data storage facility (e.g., a RAID group) and a destination storage server which has its own persistent data storage facility (e.g., another RAID group). The present invention is directed to a specific type of mirroring system that operates based on regularly occurring events called "consistency points".

Write requests received by the source storage server from clients are logged and transmitted to the destination storage server, where they are saved in a file associated with the source storage server. However, the write requests are not immediately applied to data in persistent storage by either the source or destination storage server. Instead, write requests are applied to persistent storage only at consistency points (which occur regularly), at both the source storage server and the destination storage server.

At each consistency point, all write requests accumulated since the previous consistency point are applied to persistent storage at the source storage server, and the mirror copy of data at the remote storage server is also brought into synchronization with the primary copy at the source storage server at this time; this is known as consistency point synchronization, or "CP sync". This happens at each consistency point, which may occur, for example, every few seconds during normal operation (depending on the particular system configuration).

Also, in accordance with the invention, the source storage server sends file system generated volume block number updates to the destination storage server during the synchronization phase of the consistency point. The destination storage server then uses those volume block number updates to update its mirror copy of the data via its own RAID layer. This is in contrast with certain prior systems in which the source storage server's RAID layer sent updates directly to the disks owned by the destination storage server (i.e., disk block number updates, not volume block number updates), bypassing the file system and RAID layer in the destination storage server. See generally Applicants' specification at paragraphs [0011] – [0013] and Fig. 4.

Referring now to the claims, claim 1 recites:

1. A method of operating a destination storage server to mirror a primary volume maintained by a source storage server, the method comprising:
 - receiving, at the destination storage server, a plurality of log entries from the source storage server, the plurality of log entries representing write requests received by the source storage server;
 - writing the received log entries to a file maintained by the destination storage server;
 - receiving, at the destination storage server, data from the source storage server during a synchronization phase of a consistency point, the consistency point being one of a plurality of regularly

occurring consistency points, each characterized by the saving of data specified by write requests in a set of non-volatile storage devices managed by a RAID layer in the source storage server and in a set of non-volatile storage devices managed by a RAID layer in the destination storage server, wherein said data received at the destination storage server from the source storage server during the synchronization phase of the consistency point include volume block number updates from the source storage server;

using the data received at the destination storage server from the source storage server during the synchronization phase of the consistency point, including the volume block number updates, to update a mirror volume maintained by the destination storage server, via the RAID layer in the destination storage server; and

using log entries from the file to update the mirror volume.
(Emphasis added.)

Each of Applicants' independent claims essentially recites the limitations emphasized above in bold.

The Examiner contends that Galipeau discloses all of the limitations in claim 1 except volume block number updates during synchronization. (Final Office Action, pp. 2-3). The Examiner admits that Galipeau fails to disclose volume block number updates during synchronization, but contends that Achiwa discloses such functionality. (Final Office Action, p. 3).

Applicants respectfully submit that the Examiner has mischaracterized the disclosure in Galipeau. As such, no combination of the cited references discloses all of the limitations of claim 1 or any other independent claim in the present application. Furthermore, the motivation alleged by the Examiner to combine the teachings of Galipeau and Achiwa does not withstand scrutiny.

First, the Examiner's characterization of Galipeau is inaccurate. Galipeau discloses a completely different type of mirroring approach, which is not based on

consistency points. In the system of Galipeau, a local computer system 1 stores received write requests in a Store and Forward log (col. 8, lines 4-44). The requests are then forwarded from the local computer system 1 to a remote computer system 5. At the remote system 5, the forwarded log entries are received and stored in a request log, which is subsequently used by a router process 53 to update a backup copy of data at the remote computer system 5 (col. 3, lines 21-40; col. 9, lines 16-41).

There is no disclosure or suggestion in Galipeau, however, of any regularly occurring synchronization between the times when write requests are applied to disk in the local system 1 and the times when write requests are applied to disk in the remote system 5. The system of Galipeau is not based on, and does not use, consistency points such as defined in Applicants' claims, i.e., where a consistency point is one of a plurality of regularly occurring events, each characterized by the saving of data specified by write requests in a set of non-volatile storage devices managed by a RAID layer in the source storage server and in a set of non-volatile storage devices managed by a RAID layer in the destination storage server, and during a synchronization phase of which the destination storage server receives data from the source storage server. In Galipeau there are no regularly occurring consistency points as defined in Applicants' claims, and as such, there is no disclosure of any synchronization phase of such an event, as recited in claim 1.

Galipeau does disclose that a send process 50 in the local computer 1 reads the Store and Forward Log entries periodically, e.g., every tenth of a second (col. 7, lines 56-61). However, that activity is not part of a "consistency point" as defined in Applicants' claims, since it is not tied to the action of saving data to disk (Galipeau

explicitly states that the send process 50 executes "in the background (i.e., asynchronously from other software . . .)". In addition, the router process 53 on the remote system 5 uses received log entries to update the backup copy on disk at a time which is not tied to saving data on disk at the local computer 1 (col. 9, lines 16-41). In Galipeau there is no disclosure of any explicit synchronization between the times when recently received writes are saved to disk on the local and remote computer systems 1 and 5, respectively, unless a user explicitly sends a synchronization command (col. 9, line 54-56), which is not a regularly occurring event.

Note that Applicants submitted the above arguments regarding Galipeau in their previous response, filed on 5/23/2006. However, in the present Final Office Action the Examiner makes no attempt to address those arguments, but instead simply asserts without explanation that Galipeau discloses those features. Applicants respectfully submit that the Examiner is mistaken, for the reasons discussed above. Applicants further submit that, if the current grounds of rejection are to be maintained, the Examiner is obligated to clearly and directly explain why the Examiner believes Applicants' arguments to be incorrect.

In addition, Applicants also do not find any disclosure or suggestion of the above-mentioned claim features in Achiwa either, nor has the Examiner so asserted. In particular, Achiwa is not seen to disclose or suggest, *inter alia*, receiving, at the destination storage server, data from the source storage server during a synchronization phase of a consistency point, the consistency point being one of a plurality of regularly occurring consistency points, each characterized by the saving of data specified by write requests in a set of non-volatile storage devices managed by a

RAID layer in the source storage server and in a set of non-volatile storage devices managed by a RAID layer in the destination storage server.

Therefore, no combination of the cited references discloses all of the limitations of claim 1 or any other independent claim in the present application. For at least the above reasons, therefore, the rejection is improper and should be withdrawn.

Furthermore, there must be something in the prior art to suggest the desirability of making the alleged combination, i.e., a motivation. In re Rouffet, 149 F.3d 1350, 1356 (Fed. Cir. 1998). In the Final Office Action, the Examiner alleges that the motivation would be: "an excellent coherency of volume image (paragraph 321)" (Office Action, p. 3). That is exactly the kind of unsupported, conclusory statement of motivation which the courts have held to be inadequate to support an obviousness rejection. The showing of motivation "must be clear and particular" findings of fact based on actual evidence, not merely broad conclusive statements. In re Dembiczak, 175 F.3d 994, 999 (Fed. Cir. 1999).

Paragraph 321 in Achiwa, which is the only "evidence" the Examiner cites to show the alleged motivation, is nothing more than a "boilerplate" paragraph that appears in many if not most patents. The Examiner has not shown how combining any of Achiwa's teachings with those of Galipeau would actually improve the coherency of volume images over what Galipeau already provides. Applicants respectfully submit that the alleged suggestion/motivation to combine Achiwa with Galipeau simply does not exist. For this additional reason, therefore, the rejection is improper and should be withdrawn.

Claims 7 and 13

In addition, each of independent claims 7 and 13 further recites that the write requests received by the destination storage server are stored in a file which is associated with the source storage server, or similar language. The cited art does not disclose or suggest such functionality. In Galipeau, for example, write requests are received by the remote computer system 5 and logged in a request log in the remote computer system 5 (col. 3, lines 21-40); however, there is no suggestion that the request log is in some way associated with any particular local computer system 1 or any computer system other than the remote computer system 5 in which it is maintained. Neither is Achiwa seen to disclose this feature. For this additional reason, therefore, claims 7 and 13 are further patentable over the cited art along with all of their dependent claims.

Dependent Claims

In view of the above remarks, a specific discussion of the dependent claims is considered to be unnecessary. Therefore, Applicants' silence regarding any dependent claim is not to be interpreted as agreement with, or acquiescence to, the rejection of such claim or as waiving any argument regarding that claim.

Conclusion

For the foregoing reasons, the present application is believed to be in condition for allowance, and such action is earnestly requested.

If there are any additional charges/credits, please charge/credit our deposit
account no. 02-2666.

Respectfully submitted,

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Dated: October 10, 2006


Jordan M. Becker

Reg. No. 39,602

Customer No. 48102
12400 Wilshire Blvd.
Seventh Floor
Los Angeles, CA 90025
(408) 720-8300